



Cheryl Watson

Ahead of the Performance Curve

Julie-Ann Williams



Foreword

Although it has been my real pleasure to know Cheryl Watson for more than 23 years, I didn't really know much about her. For that reason I was thrilled to learn that she was being profiled in the new book, "*Cheryl Watson – Ahead of the Performance Curve*". It's in this book where I really discovered that she is not only the sweet lady I had the privilege to see just occasionally at CMG or SHARE conferences, but that she also was a tireless worker who was always eager to learn about new technologies that came her way. Hopefully, with the exposure of this book many others will come to appreciate what a huge difference Cheryl Watson has made in improving the performance of the IBM mainframe for so many years.

It was shortly after I started publishing a magazine named *Mainframe Journal* that I met Cheryl while we were both attending a CMG (Computer Measurement Group) conference. I regularly attended CMG because it was the perfect conference to promote *Mainframe Journal* as it was almost entirely focused on the performance management and capacity planning issues facing users of IBM mainframe computer systems. The attendees were prime candidates to become subscribers and the speakers were all ideal candidates to write the technical articles I wanted to publish. And Cheryl Watson was at the top of my list! So I was thrilled when she gave me an article to publish titled, "Why Isn't a CPU Second Consistent?" That article turned out to be one of the most popular articles we ever published. From that point on, I felt honored any time we were able to feature a Cheryl Watson article.

Over the succeeding years I have witnessed the rise of Cheryl Watson to virtual mainframe rock-star status. But to me she is much more; she is a very sweet lady I am proud to call "friend". She also helped to elevate *Mainframe Journal's* reputation when she was kind enough to allow me to publish her exceptional articles.

In August 2010, I had the privilege to add Cheryl Watson to the Mainframe Hall of Fame because of the way she has shared her extensive knowledge of performance analysis and capacity planning since 1965 as a consultant, teacher, and author of "Cheryl Watson's Tuning Letter."

Bob Thomas
President, Enterprise Systems Media, Inc.
Publisher of Enterprise Management & Enterprise Tech Journal
Dallas, Texas



Performance Management

This is the second paper introducing a “Who’s Who” of some of the brightest, and yet under-sung stars of our mainframe generation. These folks have kindly agreed to be interviewed and will be providing their input to the new book, “z/Auditing Essentials VOLUME 2 – Auditing z/OS” that follows on from the acclaimed “z/Auditing Essentials VOLUME 1 - zEnterprise Hardware - An Introduction for Auditors”.

I waxed lyrical about being able to work with many of my heroes in the last paper and this time is no different. Cheryl Watson is unique in the mainframe industry. There are folks out there who have heard of “Cheryl Watson” but know it only as a corporate identity not as a person! This woman stands out from the crowd, not just because she is a woman (unusual enough on its own in this business nowadays) but also because she has helped to shape the systems that we work on since early in her career. Cheryl was one of the first people who challenged IBM’s newly issued performance figures for CMOS (**C**omplementary **M**etal **O**xide **S**emiconductor) based mainframes when they were first released. She continues to blaze an influential path today...



Cheryl Watson-Walker

As Barry Schrager was introducing the mainframe world to the idea of Data Security, Cheryl was making the dark arts associated with tuning the performance of these beasts accessible to those of us with enough understanding to have been only dangerous before!

I first came across the Cheryl Watson Tuning Letters as a nervous young MVS Systems Programmer! But more of that later, I want to start at the beginning so come with me to 1965 the year our heroine first gets involved in mainframes.

Born in Portland, Oregon (never ask a Lady how old she is!), Cheryl was something of a Nerd at High School. She’d always liked Math but accessible computers were still some way off, so the term Geek really hadn’t taken root yet!

She majored in Math at Portland State University (with a minor in Physics) and married whilst still in college. So when her husband decided to get his PhD, Cheryl had to abandon plans to get her Masters in favor of getting a job to support the two of them. Thanks to her selfless devotion to family, the mainframe world gained one of its brightest stars, although no one would have suspected it at the time!

Cheryl joined Consolidated Freightways as an Applications Programmer in 1965. They had both an IBM 1401 and the mainframe version, the 7010.

The latter had a massive four I/O Channels and 100,000 characters of memory in which to operate all instructions. It cost \$1,307,550 back in 1964!

She initially learnt to write Autocoder but, when Consolidated moved to S360 in 1966, changed to COBOL and eventually moved into the Systems Programming team.

Cheryl had found her niche! In fact, she stayed in this



You may be wondering what an interview with the woman who is most famous for questioning IBM’s mainframe performance figures is doing in a book about security...

Wikipedia say “Big Data is a term applied to data sets whose size is beyond the ability of commonly used software tools to capture, manage, and process the data within a tolerable elapsed time.”

And Big Data is exactly what we are securing and auditing on System z. Even restrained estimates still suggest that 70% of the world’s business data sits behind a mainframe. The world has at last caught up and now the true value of data is being realised.

It is equally important to make sure that our Big Data can be reached in a timely manner as it is to make sure that it is secured against misuse.

team for three years and only moved on when her husband found a new job and she went with him.

Interestingly, a pattern started to emerge. Cheryl found a job working with one of her new neighbors who was in the process of setting up a software company to produce a 4GL compiler. This wouldn't be the last time that Cheryl made career progress not because of a resumé (or CV) but because of the personal contacts she has made. In fact, even today, Cheryl has never had a regular resumé! It's a testament to her determination to succeed that she has never been asked to conform to the normal expectations set upon employees.

Exploiting her tendency to make the best of any bad situation, Cheryl turned to her advantage the simultaneous collapse of her marriage **and** the company she was working for, after a chat with an old friend. This former colleague was working for ISM (Information Systems Management) when he convinced Cheryl to move to Phoenix, Arizona to take up what would turn out to be the first of many "Dream Jobs" as a traveling Trainer for the software that ISM supported.

If there's a better way to get over a messy divorce than traveling the world being treated as an expert on many different things, I have yet to find it! So Cheryl traveled and had fun while she worked. She was teaching out in Europe (Holland to be precise) when ISM (there is now another organization using the same name) sadly went out of business!

If Cheryl wasn't such a lovely example of a human being it would be impossible not to dislike her at this point...

While she was sitting in a café (No not one of those café's) in Holland, wondering how she was going to get home, ISM was bought by an organization based out of San Francisco. They had heard of Cheryl and what she was doing for ISM and they REALLY wanted her on board. They arranged flights back to the US for her to go straight to San Francisco, packing up and moving her entire life to be there when she arrived. Now **that** Ladies and Gentlemen is some welcome!!

Quite rightly, Cheryl LOVED being in San Francisco and was there for a little over a year before she was head hunted by EDS (Electronic Data Systems founded by H. Ross Perot). The wonderful thing about having had a varied career is that you have interesting answers when people ask you questions. So a formal chat with her peers turned into yet another high profile, career enhancing job offer!

She worked for EDS (on and off) for over 10 years and had some very interesting engagements during that time. I've picked just a few examples here.

In 1971 Cheryl was placed onto the EDS team that was developing a new Claims System for Blue Cross/Blue Shield. She started on the project as a COBOL programmer but immediately demonstrated her uncanny ability to poke around under the covers of the system she was working on and unsurprisingly was swiftly moved into Systems Programming!

Whilst there, Cheryl took another one of those odd calls that she gets... Amdahl wanted her to go and work with them in California so she went. Despite only being there for 6 months she says: "I learnt **loads!**" But she soon returned to EDS in San Francisco. It was while she was there that she met and married her second husband.

They were then moved to Dallas by EDS for Cheryl to work as a Trainer in their Systems Engineering Development Scheme. This was a three month long, total immersion training plan and started on day one with machine code and Assembler. Some people questioned this sequence but Cheryl defends her position by saying: "If we started with Assembler on day one then by the time we got to COBOL, the



students could already code and more importantly, read dumps!”

Cheryl realized quite quickly that her real skills lay in a pure technical direction. Whilst she taught Management Skills, she freely admits that she never learnt the lessons for herself! However, EDS always allowed promotion through a technical career, not just by moving into management, so it was never a problem.

Cheryl and her husband returned to San Francisco to help to get a troubled project back on track. Whilst there, EDS wanted Cheryl to help with the recruitment of a full team of mainframe professionals to work for Allstate Insurance in Northbrook, Illinois. They asked her to interview candidates. She had been out to visit with the Customer three or four times when she realized that one of the jobs she was recruiting for was one that she would just love to have for herself! It was the CICS Administrator/Systems Programmer role, which involved teaching and many other things that Cheryl enjoyed. So she quite blatantly poached the job, moved to Illinois and had two years of having a surprisingly fabulous time despite her second divorce. She retook her maiden-name of Watson and prepared to launch “Cheryl Watson” on an unsuspecting public!

Cheryl stayed until the end of the project and got her first taste of Performance Management after which EDS moved her back to Dallas in 1978.

RMF (Resource Measurement Facility) was in its infancy but when Cheryl saw its predecessor, MF1 (Measurement Facility 1) she just thought it was the neatest thing she'd ever seen! She eventually did a deal to exchange a large amount of unused vacation time for attendance at her first CMG Conference (Computer Measurement Group). User Groups can have a dramatic impact on people when they first discover them and Cheryl was not immune. She reports being “completely blown away” after hearing Barry Merrill talking about one of the first Performance Monitoring tools, MXG (Merrill's Expanded Guide which was, and still is, based on SAS).



Cheryl wrote up a 100 page report on her attendance at CMG which she shared with all of her colleagues. She started with a cover page summary, followed by 2-3 pages titled “Things We Should Be Doing”. Then each department had its own section, which incorporated both a summary of the relevant presentations and a copy of the handouts provided at the conference. She also gave classes based on information she had learned at the conference!

It might seem like a lot of effort to go to but these actions guaranteed that Cheryl's name was top of the list when her employers were looking to send someone to a conference after that. She started volunteering with various User Groups in the late 1970s.

Cheryl takes up the story a while later:

“One of the best times I had working came in 1980s when EDS sent me to the Netherlands to work on an outsourcing assignment. I was able to do all of the things I loved, such as teaching, performance, and capacity planning, as well as attend some of the European conferences.

That's where I re-connected with John McCann, who I had worked with at SHARE in the CME (Computer Measurement and Evaluation) project. That was a happy meeting, because a year later, John hired me to help start up the German office of Morino Associates (founded by Mario Morino). That was almost two years of living half-time in the UK and half-time in Germany, and traveling throughout Europe teaching the Morino

The logo for AMDAHL, consisting of the word 'amdahl' in a lowercase, red, sans-serif font.

products and providing technical sales support.


It was a dream job, and let me see so much of Europe and develop some great friendships. That's also where I learned SAS for the first time."

Cheryl returned to the USA when Morino (later merged with another company to form Legent Corporation in 1989 and subsequently taken over by CA in 1995) wanted her for their Education Department. And so she started teaching products like TSO/MON (a TSO performance manager product) and MICS (MVS Integrated Control System - a SAS database product now marketed by CA). Cheryl is fiercely loyal and I love the way she always suffixes TSO/MON with the statement "Still one of the greatest products I've ever seen!"

This high profile job saw Cheryl meeting with hundreds of people in the Performance Management and Capacity Planning fields. But after a while the call of the deeply technical became unavoidable and she moved into development of a major release of MICS. It seems that Mario Morino was a very influential person in Cheryl's life. She says: "His passion for the field was infectious!"

So infectious it seems that Cheryl was soon working 50-60 hour weeks at Morino Associates Inc!

It was during this particularly intense period that Cheryl met the man who turned out to truly be her soul-mate, Tom Walker. Cheryl is unapologetic about championing her relationship with Tom and she has good reason for that. But I'll let her tell the next part of the story:



Tom Walker

Tom Walker's point of view:

"Deciding to start and finance a company based on Cheryl's talents was easy after seeing the rave reviews from her students and readers. Cheryl's enthusiasm is contagious, in person and in writing.

The natural next step was to overrule her natural modesty and focus all our marketing and our products on her experience and personality.

As an editor and having mainframe but no IBM experience, I happily functioned as the most clueless reader, asking the dumb questions and getting Cheryl to expand her explanations to accommodate people new to the subject.

Working together 24/7 since 1986 might sound difficult, but it's not for us. We're very compatible, laugh all the time, and even had identical furniture when we met!"

"Tom had just retired from working for three years for ITT, who had bought a time-sharing company he had part ownership in (Dialcomm, Inc., a forerunner of companies such as AOL). He had been in mainframe computers since 1965, but not IBM (GE, Honeywell, etc.). Since I couldn't work those hours and find time to spend with Tom, I resigned and went to work as an independent consultant and trainer."

Whilst working as an independent, Cheryl started writing articles on performance for Mainframe Journal and Technical Support Magazine. These became very popular, especially the multi-part articles on RMF and another set on SRM (Storage Resource Management). Back to Cheryl:

"In 1986, Tom decided to finance our joint company and use my name as the major item in advertising our classes.

I developed the classes, Tom reviewed and edited them, and he financed the startup. That startup included huge mailings of class advertisements, all with my name in large bold letters on the ads.

It worked, and the courses took off. We taught throughout the world - throughout the US, England, Europe, South Africa, Scandinavia, Singapore, etc. until we wanted to stop traveling. Then we simply taught classes in Florida, where we moved in 1989."

As we all know, none of us are getting any younger, and in 1991 Cheryl and Tom took the decision to change the direction of the company. Instead of her exhausting schedule traveling around the world giving classes, Cheryl began to write "Cheryl Watson's Tuning Letter". This started out as a 16 page monthly (which grew to 40-50 pages bi-monthly) newsletter devoted to: How to tune the mainframe and How to do Capacity Planning.

With Tom's help as Editor, Cheryl quickly learnt how to write well. His influence was visible in the marketing sphere too (he broke with "tradition" sending the whole of a

newsletter as a flier instead of just a teaser which was normal practice at the time). The accessible style that this pair has developed has been an inspiration to my own writing style. Cheryl assures me this is Tom's influence at work and that her natural style is firmly embedded in Geek!

In 1995 Cheryl began pushing WLM (Work Load Manager) into our collective consciousnesses. She immediately realized that this was the future and that performance to specified targets could only be a good thing. Sadly we in the User Community mostly disagreed! My own reason for not liking WLM in the early days was firmly based in the fact that, with Cheryl's help, I had finally got to grips with IPS and ICS in SYS1.PARMLIB (**I**nstallation **P**erformance **S**pecifications and **I**nstallation **C**ontrol **S**pecifications) and didn't want to accept that WLM meant my hard earned knowledge was now pretty much useless!

Cheryl's reaction to this wasn't to run away and hide in a corner (as I have been known to do if things don't go the way I want – I admit it, sometimes I revert to being a 6 year old!). Oh no! Her reaction was to work with IBM to get the length of time we had to convert to WLM extended and to publish a starter set of WLM policies known as QuickStart. It has been this patient, hand-holding rather than brow-beating approach that is Cheryl's trademark and part of why she still remains so popular and relevant today.

One of the other reasons that Cheryl has maintained a highly relevant technical profile in the mainframe industry is that she is always looking for a new cause to champion. Her next target after convincing us all that WLM was achievable was to demystify the archaic performance charts that IBM uses to rate their mainframe systems.

It took a great deal of study to figure out the process used by IBM and, as soon as she had done so, it was time to share with the wider community. Her first port of call was to publish her latest newsletter with details on how LSPR (**L**arge **S**ystems **P**erformance **R**eference) and processor speeds were analyzed.

The landscape for capacity modeling effectively changed overnight with the switch to CMOS based machines. Instead of modeling a single large processor we suddenly had to deal with multiple smaller processors which could work together.

Cheryl took the campaign slightly further than just telling people what she knew. This time she designed and wrote a product, BoxScore, as Cheryl put it: "To make sure that you are getting what you are paying for!"


BoxScore is still a vital tool for many organizations when it comes to modeling changes. It can also identify some performance type problems e.g. some jobs run **WAY LONGER** than others performing similar function.

Cheryl identified that this can sometimes be due to poor coding standards, for example, using subscripts instead of indexing when writing COBOL. Increasing your system capacity would make no difference to the performance of such examples of "bad coding".

In 1998 Cheryl suffered a heart attack after a horrific home invasion. She exercised herself back to 80% fitness within a year! On top of that, in 1999, she was awarded

BoxScore

BoxScore is an informative suite of programs, developed by Cheryl Watson that shows you the difference in CPU speed and CPU capacity after a hardware change, a software change, a change in LPAR configuration, or a tuning change.



the CMG Michelson Award. In 2000, she and Tom married, and she became 'Cheryl Watson Walker', using Walker at home and Watson in business. I'm not sure that there is anything that can keep this Lady down for long!

Cheryl has always been passionate about the things that she believes in. Her latest campaign is based in how excited she is about z/OSMF! So what is z/OSMF?

IBM says:

"The z/OS Management Facility is intended to enable system programmers to more easily manage and administer a mainframe system by simplifying day to day operations and administration of a z/OS system. "

Cheryl Watson says:

"z/OSMF IS the future for z/OS administration!"

She acknowledges that mainframe specialists don't currently "trust" the facility but that this is because IBM made some odd decisions about its direction in the early days. Cheryl has been fighting our corner with IBM to make sure that the focus is on making the facility run better before pushing out new functionality.

She sees the slow take up of z/OSMF as being caused by the same problem as was faced by WLM in the early days. Overworked Systems Programmers don't see the immediate benefit to them. However, we have to start thinking about the long-term future for **my** platform of choice, the IBM mainframe.

Lack of investment in new mainframe staff over the last few decades means that we have to take creative approaches to up skilling the Admins of the Future. And by Admins we are going to have to start envisioning people performing TCP/IP configuration, WLM changes, Capacity Management, software distribution, dump analysis and other, more classic ISPF (**I**nteractive **S**ystem **P**roductivity **F**acility) tasks. The wizards built into z/OSMF can simplify the tasks and reduce the learning time for new personnel.

Cheryl has kindly agreed to provide us with her personal insights into z/OSMF and the future of mainframe administration. But before we get to that, I'd like to share an extract from Cheryl Watson's first ever Tuning Letter in Appendix 1. When we were talking about which piece to include here, Cheryl and I both wanted to give the reader a feeling of the gravitas of operating in an environment that has been around since the 60s.

Cheryl explains it best: "Because the basics of MVS, the under-pinnings if you will, have not greatly changed, most of the old Tuning Letters contain articles that are relevant in today's environment (and are still available). These two articles from my first Tuning Letter in January 1991 illustrate that nicely. The MVS Review relating to program search is still the same, with the slight addition of a new PROGxx member. The Reducing ISPF Response Time is still valid, but less important because the default IBM distribution now uses the preferred placement of ISPF LPA (**L**ink **P**ack **A**rea) modules."

The extract is available in Appendix 1 whilst the full Tuning Letter can also be found here www.watsonwalker.com/JAN91.pdf.

So where does this leave us? We have more and more mainframes being used, but are left with a small, and shrinking, number of staff to look after them. Ooops!

But don't panic! IBM is coming to our aid with z/OSMF, making the everyday tasks (and some of the not so every day) more accessible to other staff. We can concentrate our skills on the more technical work. z/OSMF can help to free up our time

so that we can get our organization's 'Big Data' stored securely, whilst maintaining high availability, but still passing those all-important audits. This is what we all want after all. Appendix 2 allows Cheryl to go into more detail about what z/OSMF is and how it can help you.

Anything that allows us to improve the performance of our platform is important. It means we can get to work on those tasks that have been sitting on the back burner and are still waiting now. We may now have the chance and “z/Auditing Essentials VOLUME 2 – Auditing z/OS” will help you!

I hope you enjoy Cheryl's new paper (in Appendix 2) as much as I did.



Remember: Civilization runs on MVS – Julie-Ann
(Shamelessly adapted from one of Bob Rogers' old tag lines)

Roll call of my Mainframe Heroes

There is not enough space here to give credit to all of the individuals who have made such a difference to the industry I work in but this should give you enough to be able to Google them yourself.

In alphabetical order...

<i>Barry Merrill</i>	<i>who... was a pioneer in performance management</i>
<i>Barry Schragar</i>	<i>who... introduced us to Data Security</i>
<i>Ben Riggins</i>	<i>who... gave us CICS</i>
<i>Bob Rogers</i>	<i>who... gave us ASM in MVS and then 64bit in z/OS amongst other things</i>
<i>Cheryl Watson</i>	<i>who... introduced us to Performance and Capacity Management</i>
<i>Eldon Worley</i>	<i>who... gave us RACF</i>
<i>Gene Amdahl</i>	<i>who... was the Chief Architect for S/360</i>
<i>Karl-Heinz Strassemeyer</i>	<i>who... brought Linux to System z</i>
<i>Kenneth Kolence</i>	<i>who... co-founded Boole and Babbage</i>
<i>Mike Cowlshaw</i>	<i>who... brought Rexx into the world</i>
<i>Pete Clarke</i>	<i>who... saved VSE!</i>
<i>Peter Relson</i>	<i>who... brought us the IBM Health Checker for z/OS</i>

Check out the “Mainframe Hall of Fame” for more influential individuals who have helped to shape my platform of choice.

<http://www.mainframezone.com/mainframe-hall-of-fame>



TUNING Letter

A PRACTICAL JOURNAL OF z/OS TUNING AND MEASUREMENT ADVICE

Editor's note: This is an extract from our January 1991 issue, which was the first newsletter we ever published. Back then, it was a monthly newsletter. Today the TUNING Letter is published six times a year. Although they vary in size, the average issue is now 44 pages.

WELCOME!

Welcome to the first issue of **TUNING LETTER!** This monthly newsletter is a natural outgrowth of my MVS classes, articles, and over twenty four years in IBM mainframe software and tuning. People *do* want to know how to tune MVS. Not everyone (especially me) concurs with vendors who say "It's simply too expensive to tune. It's cheaper to upgrade and buy more hardware". In reality, it's seldom cheaper to buy more hardware!

This newsletter is written for the systems programmer or performance analyst who wants to **1) extend the life of their current system, 2) provide better service to their users and/or 3) provide room for growth without the cost of upgrading their hardware.**

Or, if they've already ordered an upgrade, it will show how to tune their systems to last until the upgrade is installed. Each tuning tip will provide one of these four benefits, saving money and/or increasing system responsiveness. Capacity planners and charge back analysts will benefit from the emphasis on MVS measurement theory and practice.

Each issue will focus on one or two aspects of MVS by providing several practical techniques to improve performance. For example, this month's issue provides seven different techniques for tuning TSO subsystems. Most articles will be applicable to both MVS/XA and MVS/ESA systems, and I will indicate if there's a difference between the two.

Each tuning tip is described in multiple sections: 1) an overview describes the intent of the change and its applicability. 2) WHAT TO DO is a step-by-step checklist of how to make the change. 3) BENEFITS describes the positive sides of the change. 4) COSTS tells you what resource(s) you're trading off! (Very few items come free: for example, you might give up some storage to reduce CPU time or give up a little CPU time to reduce I/Os). 5) HOW TO MEASURE provides techniques for determining whether you've made an improvement. Naturally, you'll always want to measure before and after a change to gauge its effectiveness for you. All in all, this will be a very practical guide to tuning your system and saving money.

THIS ISSUE

FOCUS: TUNING TSO

Reducing ISPF Response Time	2
Eliminating TSO STEPLIBs	4
Reducing Output Wait Swaps	6
Reducing TSO Period Swaps	9
Invoking CLISTs Faster	12
Reducing CLIST Response Time	13
Providing Consistent Response	14

MVS REVIEW

MVS Program Fetch	16
-------------------	----

MVS MEASUREMENTS

Measuring TSO Response Time	17
Measuring TSO Command Use	18

Q & A	19
-------	----

MISCELLANY	19
------------	----

FOCUS: TUNING TSO

Most installations see at least a 30% increase in TSO CPU use each year. This rate of growth appears to be continuing or even increasing as more TSO applications and 4GL products are installed. Tuning TSO subsystems can often provide a reduction in TSO response time of 50% or more.

This month's focus provides seven ways to improve or adjust TSO response time and reduce its resource consumption. The impact of any change will depend on the installation. The techniques are described in descending order of potential savings for a typical TSO site.

Reducing ISPF Response Time

[Editor's Note - The following tuning tip is usually not needed for ISPF in most installations today because IBM distributes its server paks with ISPF modules in LPA. If you've overridden IBM's default LPA assignments, this article is still applicable to you. Additionally, any references to expanded storage can be ignored. If you are considering using LPA for another application, the same considerations will apply.]

You can reduce ISPF response time, average TSO swap size, and central storage use by moving ISPF modules to LPA (Link Pack Area) from their default location in linklist. The majority of TSO commands are ISPF commands, and would be found much faster in LPA than in link list. See MVS PROGRAM FETCH in this newsletter. Consequently, moving ISPF modules to LPA reduces response time.

In addition, modules found in linklist are loaded into the user private region and are liable for swap out and swap in. Because they're in the private region and not common storage, you might actually have multiple copies of the same module in central storage. Consequently, leaving ISPF modules in linklist results in additional CPU time and I/O for loading, and larger working set sizes (WSS) for TSO users. Larger WSS result in larger swap sets and therefore more pages are moved during a swap. If logical swap is in effect, this larger WSS uses more central storage. Even if logical swap is not used, the resulting physical swap to auxiliary storage requires more I/O. Hence, moving ISPF modules to PLPA is an excellent tuning step for any ISPF shop.

WHAT TO DO

The solution is simple - place reentrant ISPF modules into link pack area. Here's a recommended checklist:

1. Test the change (under an LPAR or stand-alone environment):
 - Measure TSO response time, average swap size, LPA size in virtual storage, and activity on the load ISPF libraries. See the HOW TO MEASURE below.
 - Make a backup of members LPALSTxx and LNKLSTxx. *[Editor's Note - Today we recommend that these changes be made in member PROGxx.]*
 - Place all reentrant ISPF modules into a separate library.
 - Add the library name to LPALSTxx.
 - IPL with CLPA (Clear Link Pack Area). This causes the new LPALSTxx to be used during IPL.
 - Measure the same indicators as above.

2. Install in production:
 - Move the changes to the production system.
 - Remove the reentrant modules from the linklist library (or simply remove the library from the LNKLSTxx member).

BENEFITS

This change will result in shorter response times, less CPU time, fewer I/Os to linklist, less use of central storage, and a smaller TSO swap size.

COSTS

This change will require more virtual storage, primarily above the 16MB line, although some modules might be below the 16MB line. It requires an IPL. It will slightly increase LPA search time for all other users since the LPA in-core directory is slightly longer. (This is minimal).

HOW TO MEASURE

This is a tough one to measure since measurements will have to be done before and after an IPL and the differing workloads may alter the results. You should analyze TSO response time before and after the change. See MEASURING TSO RESPONSE TIME at the end of this newsletter. [Editor's Note - you can find the entire newsletter at www.watsonwalker.com/sampleissues.html.] This could be done for one TSO user issuing a variety of ISPF commands or by simply looking at the average response time for all TSO users. You'll also want to look at virtual storage usage, swap sizes and activity on the ISPF libraries. Figure 1 shows a portion of an RMF Virtual Storage Activity

STATIC STORAGE MAP		
AREA	ADDRESS	SIZE
EPVT	2E00000	2002M
ECSA	1EB5000	15.3M
EMLPA	1EAA000	44K
EFLPA	0	0K
EPLPA	1AB8000	4040K
ESQA	1268000	8512K
ENUC	1000000	2462K
----- 16 MEG BOUNDARY -----		
NUCLEUS	F96000	423K
SQA	EA6000	960K
PLPA	CEC000	1768K
FLPA	0	0K
MLPA	CEA000	8K
CSA	600000	7080K
PRIVATE	1000	6140K
PSA	0	4K

Figure 1 - RMF Virtual Storage Activity Report

WORKLOAD ACTIVITY												
OS/VS2 SP2.2.0		SYSTEM ID SYSA RPT VERSION 3.5.0		START 08/16/89-10.00.00 END 08/16/89-11.00.00		INTERVAL 001.00.00 IPS = IEAIPS00		SU/SEC= 700.0 SRB = 10.0				
OPT = IEAOPT00 ICS = IEAICS00		REPORT BY PERFORMANCE GROUP PERIOD				SERVICE DEFINITION COEFFICIENTS IOC = 5.0 CPU = 10.0 MSO = 3.0				SU/SEC= 700.0 SRB = 10.0		
GROUP NUMBER	GROUP PERIOD	OBJTVE NUMBER	DOMAIN	TIME SLICE GROUP	INTERVAL (TOTAL, BY TYPE AND PER SECOND)	SERVICE	AVERAGE ABSORPTION, AVG TRX SERV RATE, WORKLOAD LEVEL	PAGE IN RATE	STORAGE AVERAGE, TOTAL	AVERAGE TRANS, MPL	ENDED TRANS, #SWAPS	AVG TRANS TIME, STD DEV HHH.MM.SS.TTT
0002	1	11	003	**	IOC= 605,069 CPU= 4,292,360 MSO= 185,918 SRB= 639,476 TOT= 5,722,823 PER SEC= 1,590	ABSRPTN = 1,445 TRX SERV= 1,247 WKLD LEV= 3.21	5.80	44856. 57968.	1.49 1.29	18230 18063	000.00.00.156 000.00.01.720	
0002	2	11	023	**	IOC= 1,209,703 CPU= 5,833,925 MSO= 230,108 SRB= 334,691 TOT= 7,608,427 PER SEC= 2,113	ABSRPTN = 1,433 TRX SERV= 1,419 WKLD LEV= .98	2.94	8501.0 12538.	1.48 1.47	3148 3366	000.00.01.235 000.00.02.277	
0002	3	11	004	**	IOC= 7,949,566 CPU= 42.190407M MSO= 2,727,881 SRB= 1,569,228 TOT= 54.437082M PER SEC= 15,121	ABSRPTN = 2,370 TRX SERV= 2,286 WKLD LEV= 64.78	.52	253.32 1616.8	6.61 6.38	2269 5064	000.00.12.480 000.00.56.810	
0002	ALL	ALL	ALL	ALL	IOC= 9,764,338 CPU= 53.316692M MSO= 2,143,907 SRB= 2,543,395 TOT= 67.768332M PER SEC= 18,825	ABSRPTN = 2,088 TRX SERV= 1,989 WKLD LEV= 51.71	1.65	7882.5 72123.	9.60 9.14	23647 26493	000.00.01.482 000.00.18.040	

Figure 2 - RMF Workload Activity Report

F														MIG=319 CPU= 97 UIC= 38 PDT= 63 DPR= 39 ASD T													
09:03:41														P P C R DP RS ESF ESF TAR WS TX WRK CPU I/O STM													
JOBNAME DMN G P L LS PR F +RS WSS IN SC RV RV RV RV																											
TSOUS13	3	2	1	WT	TI	FF	0	234	234	0	31	0	15	+0	+0	+0											
TSOUS27	4	2	3	IN		68	115	141	256	0	47	0	10	+0	+0	+0											
TSOUS05	4	2	3	IN		62	95	406	501	0	38	0	10	+0	+0	+0											

Figure 3 - RMF Monitor II ASD Screen

DEV	VOLUME	LCU	DEVICE	AVG	AVG	AVG	AVG	AVG	AVG	AVG	...	AVG
NUM	SERIAL		ACTIVITY	RESP	IOSQ	CUB	DB	PEND	DISC	CONN	...	DS
			RATE	TIME	TIME	DELAY	DELAY	TIME	TIME	TIME	...	OPEN
100	DSK010	00D	0.013	11	0	0.0	0.1	0.1	8.9	1.9	...	2.6
101	DSK102	00D	4.200	18	2	0.0	0.1	0.1	12.2	3.7	...	23.4
102	DSK203	00D	21.507	73	39	0.0	0.1	0.2	25.6	7.6	...	124
103	DSK024	00D	23.146	85	51	0.0	0.1	0.2	25.3	8.6	...	124

Figure 4 - RMF Monitor I DASD Device Activity Report

report. Before and after the change, look at the amount of LPA above (EPLPA) and below (PLPA) the line, and the PRIVATE region below the line. You can expect to see the EPLPA and PLPA increase, while the PRIVATE area might decrease.

The average swap size can be found in two places. The average swap size for all swaps (TSO as well as non-TSO) can be found in the RMF Monitor I Swap Placement Activity report shown in Figure 5. The second line from the bottom shows the average swap size, "AVERAGE PAGES PER SWAP IN - 102". Moving ISPF modules to LPA should reduce this number, particularly where the largest percent of swaps is for TSO.

The RMF Monitor I Workload Activity report in Figure 2 shows the average number of central and expanded storage frames for users within a performance group (under the column STORAGE AVERAGE). This figure should also decrease. You can get detailed information on each TSO user with RMF Monitor II data (or any other online monitor) as shown in Figure 3. The ASD (Address Space State data) includes the WS IN field which shows the average working set size (in pages) for swapping. Implementation of this change should reduce both the WS IN, as well as the number of real storage frames (RS F) and expanded storage frames (ESF). In MVS/ESA, real storage is referred to as central storage and the column is titled CS F.

Don't forget to look at the device activity to the ISPF load library devices (RMF type 74 - see Figure 4). You will see a decrease in the activity rate (SSCHs per second) and the average DASD response time for the volume(s).

This sounds like a lot of trouble, but it's only looking at four reports and you should always know the impact of any change. You'll definitely see improvement in at least one of them and probably more.


```

OS/VS2                SYSTEM ID SYSA                START 08/16/88-08.14.02  INTERVAL 007.59.50
SP2.2.0              RPT VERSION 3.5.0                END   08/16/88-15.59.01  CYCLE 1.000 SECONDS

OPT = IEAOPT01                S W A P   P L A C E M E N T   A C T I V I T Y

                                *----- AUX STORAGE -----* *---LOGICAL SWAP---* *--EXPANDED STORAGE--*

                                AUX
                                STOR
                                TOTAL
                                AUX STOR
                                DIRECT
                                AUX STOR
                                VIA
                                TRANSITION
                                LOG SWAP
                                LOG SWAP
                                EFFECTIVE

TERMINAL      CT   381,599  13,011   7,518   5,493  366,584  326,598
INPUT/OUTPUT RT    13.25    0.45    0.26    0.19   12.73    11.34
WAIT         %    86.0%    3.4%   57.8%   42.2%   96.1%    89.1%

LONG                   CT    7,194    18      2      16     6,498    5,432
WAIT                   RT    0.25    0.00    0.00    0.00    0.23     0.19
                       %    1.6%    0.3%   11.1%   88.9%   90.3%    83.6%

DETECTED               CT   10,726    420     22     398   10,058    7,572
WAIT                   RT    0.37    0.01    0.00    0.01    0.35     0.26
                       %    2.4%    3.9%   5.2%   94.8%   93.8%    75.3%

UNILATERAL   CT   41,958    370     142     228   38,541    34,013
                       RT    1.46    0.01    0.00    0.01    1.34     1.18
                       %    9.5%    0.9%   38.4%   61.6%   91.9%    88.3%

EXCHANGE ON   CT    1,975    18      1      17     1,875    1,492
RECOMMENDA-  RT    0.07    0.00    0.00    0.00    0.07     0.05
TION VALUE   %    0.4%    0.9%   5.6%   94.4%   94.9%    79.6%

ENQUEUE               CT    275     4      3      1     243     201
EXCHANGE               RT    0.01    0.00    0.00    0.00    0.01     0.01
                       %    0.1%    1.5%   75.0%   25.0%   88.4%    82.7%

TRANSITION             CT    120     119    119     0      1        1
TO NON-                RT    0.00    0.00    0.00    0.00    0.00     0.00
SWAPPABLE              %    0.0%   99.2%  100.0%   0.0%    0.8%   100.0%

TOTAL                  CT  443,847  13,960  7,807  6,153  423,800  375,309
                       RT   15.42    0.48    0.27    0.21   14.72    13.04
                       %  100.0%   3.1%   55.9%   44.1%   95.5%    88.6%

AUXILIARY STORAGE - AVERAGE PAGES PER SWAP OUT - 65   AVERAGE PAGES PER SWAP IN - 102
OCCURRANCES OF TERMINAL OUTPUT WAIT - 35,703

```

Figure 5 - RMF Monitor I Swap Placement Activity

MVS REVIEW

This section is provided to explain basic MVS concepts or techniques that relate to material presented in the tuning or measurement articles. This month, a look at how the MVS Program Fetch operates.

MVS Program Fetch

MVS program fetch locates and loads programs and TSO commands and is used for program loads, links, XCTLs, attaches, initial executes, TSO commands and CLISTS.

Knowledge of its search order will help you understand several tuning tips. Figure 6 shows the search order.

Step 1, the job pack area, refers to programs that have been previously loaded and are still resident in the address space. It will seldom find modules in this step.

Step 2 refers to the dataset (TASKLIB) that is defined during the ATTACH of a task. Applications seldom use this technique, but it's often used by software products, such as CICS with the DFHRPL ddname. TASKLIBs are product dependent and their installation instructions will tell you when to include them in the JCL. Modules from TASKLIB are loaded into the user's private region and increase his working set size (WSS).

Step 3 is to search STEPLIB and all of its concatenations. Batch applications use STEPLIBs to point to test load libraries that are continually updated. Some sites put STEPLIBs in TSO LOGON procs which point to ISPF load libraries or production load libraries. Sysprogs might use TSO STEPLIBs for testing new releases of software products. Remember that STEPLIB is searched for every TSO command. Modules from STEPLIB are loaded into the user's private region and increase his WSS.

Step 4 is to search JOBLIB, but only if there isn't a STEPLIB. The datasets concatenated under the JOBLIB ddname are searched. Started tasks and TSO users never have JOBLIBs. Modules from JOBLIB are loaded into the user's private region and increase his WSS.

Step 5 is a fairly short step. Member LPALSTxx in SYS1.PARMLIB defines one or more libraries containing reentrant modules that are to be loaded at IPL into common virtual storage (called LPA - Link Pack Area). A directory of all LPA modules is also created at that time. As modules are referenced, a pointer is added to the "active CDE (Contents Directory Entry)" list. Some of the LPA modules may be defined in SYS1.PARMLIB member IEAFIXxx as residing in fixed storage (FLPA), while other LPA modules may be defined in IEALPAXx as being modifiable (MLPA). Step 5 is a search through the FLPA, the MLPA and the active CDE list. Most active routines, such as access method routines and high use TSO routines will be found in this list. *[Editor's Note - In today's z/OS, dynamic LPA modules, specified in PROGxx members, are searched first before FLPA.]*

- 1. Job pack area**
- 2. TASKLIB(S)**
- 3. STEPLIB if present**
- 4. JOBLIB if present and no STEPLIB**
- 5. Link pack area - active modules**
- 6. Link pack directory - inactive modules**
- 7. SYS1.LINKLIB and LNKLST modules**
- 8. If TSO, SYSPROC libraries**

Figure 6 - Program Fetch Search

Step 6 goes on to search the rest of the link pack directory, an in-core table that is dependent on the number of modules in SYS1.LPALIB and datasets defined in LPA.

Step 7 searches the linklist members. This list consists of SYS1.LINKLIB and all datasets defined in LNKSTxx in SYS1.PARMLIB. Most installations will have all or part of the linklist directory in storage using LLA (Library Lookaside). Modules from linklist are loaded into the user's private region and increase his WSS. *[Editor's Note - z/OS allows you to change the order of this search with the PROGxx member.]*

Step 8 (finally!). If program fetch has not found the module, it assumes that it's searching for a TSO CLIST and starts a search of the datasets concatenated under the SYSPROC DD statement of the TSO logon proc. This is where you would expect to find TSO CLISTS.

You'll see in steps 2, 3, 4 and 8 that a physical DASD search of the library directories must be made. The physical directory searches can be fairly costly in response time and also in the amount of contention on a DASD volume. The elimination of the DASD directory search time is one of the benefits of using LPA for modules. Another major benefit is that modules in LPA reside in common storage while other modules reside in the private address space and take more central storage and result in larger swap sizes.

In MVS/ESA, there is an additional aspect of library searches to consider. Any library can be controlled using LLA (Library Lookaside). This includes TASKLIBs, STEPLIBs, JOBLIBs, and of course linklist. If a library is controlled by LLA, program fetch will request the directory search of LLA. This results in a cross memory search (using CPU cycles rather than I/O to search the directory). Paging may be involved in place of traditional I/O if the directory is not in central storage at the time of the request. Installations with a lot of expanded storage will tend to find most directories living in central storage or expanded storage. Use of LLA reduces the physical directory search, although it does not reduce the working set size or additional use of central storage.

[Editor's Note - The MVS Initialization & Tuning Guide, SA22-7591, now describes this search order in Chapter 1, 'Placing Modules in the System's Search Order for Programs.]

SPECIAL ARTICLE

Cheryl Watson's *Tuning Letter*

See below for subscription information.

z/OSMF



This special article has been extracted from several articles that appeared in Cheryl Watson's *Tuning Letters* in 2010 and 2011.

System z Simplification

z/OSMF is the latest part of a \$100 million five-year effort to simplify the System z platform. The first results of this effort (described in an October 4, 2006 [press release](#)) were announced with z/OS 1.8 and included facilities like the IBM Health Checker for z/OS (Tuning Letter 2010 No. 2, page 27), the OMEGAMON z/OS Management Console (Tuning Letter 2006 No. 2, page 45), Hardware Configuration Manager (HCM) enhancements (Tuning Letter 2006 No. 4, page 30), IBM Configuration Assistant for z/OS Communications Server (Tuning Letter 2010, page 5), application development simplification, Predictive Failure Analysis (PFA) (Tuning Letter 2010 No. 2, page 28), and the z/OS Basic Skills Information Center (Tuning Letter 2010 No. 2, page 15). The result provides a z/OS platform that takes fewer people to manage, and ensures that z/OS continues to provide the best ROI (Return on Investment) in the industry.

z/OSMF is
one of the most
exciting products
I've seen since WLM.
And it's free!

Since that time, these facilities have been expanded and new facilities have been added. There is no charge for any of these new facilities, but you need to install them and configure them. Not enough installations have activated these, and it's a shame because they will save sysprog and operator time after installation.

In this special article, we describe the latest entry into this group of simplification tools - z/OSMF. z/OSMF (z/OS Management Facility) is a Web browser-based tool to let a system programmer install and maintain a sysplex in less time using more automated tasks. z/OSMF is the most exciting z/OS enhancement since Workload Manager! In addition, this tool is free, and many of the functions can be run on a zIIP or zAAP if you have one available.

This document is reprinted with permission from **Cheryl Watson's Tuning Letter**.
Entire content copyright © 2012. Watson & Walker, Inc., 7618 Sandalwood Way, Sarasota, FL 34231.
Cheryl Watson's Tuning Letter is a practical journal of z/OS tuning and measurement
published six times a year since 1991. All rights reserved.
2012 SUBSCRIPTION RATES: Electronic version (email) \$1050 per year includes DVD of issues since 1991.
Multi-site discounts available. Subscribe online at www.watsonwalker.com.

Introduction

What is z/OSMF?

z/OSMF is the new, no-charge, feature that IBM rolled out with z/OS 1.11 (in September 2009) to help z/OS system programmers. Most z/OSMF 1.11 functions are also available in z/OS 1.10, and can run on z/OS 1.12 and 1.13. z/OSMF 1.12 can only run on z/OS 1.12 and 1.13, while z/OSMF 1.13 only runs on z/OS 1.13.

The more I see and use z/OSMF, the more I'm impressed with it. This is the most important new facility since Workload Manager and Parallel Sysplex. I normally don't get behind a product right away unless I believe that IBM is putting enough resources behind it and that the user community needs the product. I think the community is overdue for this product, and I know that IBM is putting resources behind it. The future for the IBM z/OS system programmers is z/OSMF. This is NOT like a few other products that fell along the way, like eWLM, z/OS Software Manager, MSYS for setup, and zMC.

Your sysprogs
WILL be using
z/OSMF
in the future!

Yet many sysprogs are reluctant to implement z/OSMF because the first release was difficult to install and required a large amount of resources. z/OSMF 1.12 and 1.13 greatly reduced both of these concerns. And there are more improvements on the way.

The reason for my continued support of z/OSMF is that I believe z/OSMF will be a requirement in future releases, not just an option. One example from IBM's Statements of Direction (SODs): the TCP/IP downloadable Configuration Assistant won't be available after z/OS 1.13. The only alternative for this function will be to use z/OSMF.

So here's an introduction to z/OSMF that provides the latest status, references, resource requirements, enhancements to z/OSMF 1.12 and 1.13, and our recommendations for getting started on z/OSMF.

z/OSMF has been presented at several conferences (many thanks to **Anuja Deedwaniya** and **Greg Daynes**, both of IBM). Based on early presentations, we expect that the following sysprog tasks will eventually be addressed under the scope of z/OSMF:

- Trouble-shooting and problem determination
- Installation and migration
- Preventive or corrective maintenance
- Configuration of new facilities or products
- Modification of current facilities
- Enablement of new functions

This is a big deal! Not everything is available in the first release, but certainly the framework for all of these is present.

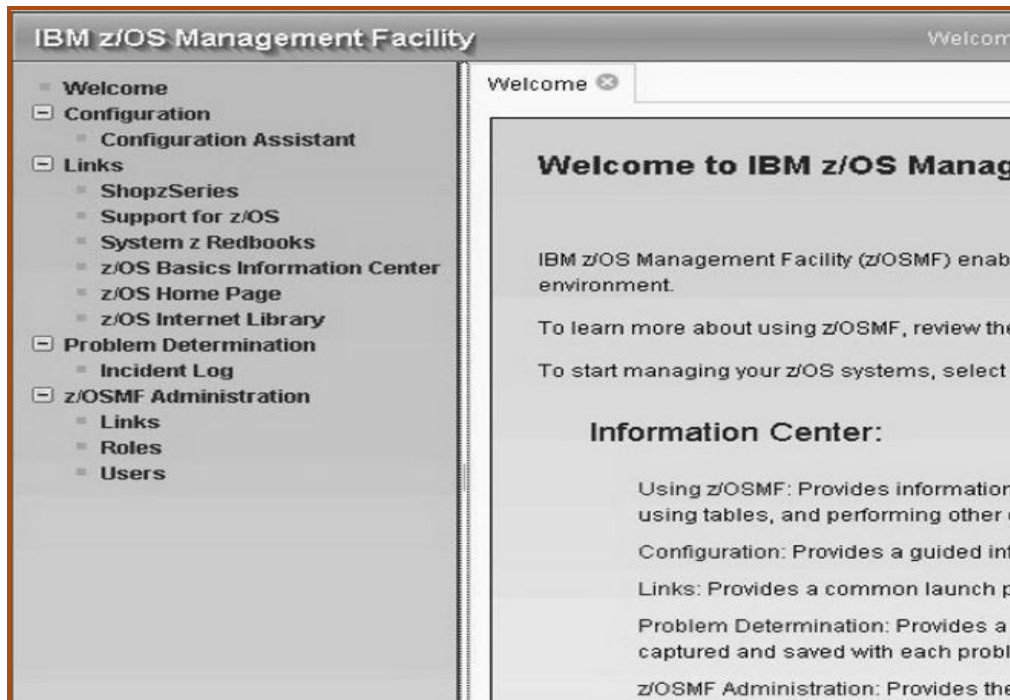
I don't intend to duplicate the manual for z/OSMF (see the *References* below). But I do want to share my enthusiasm for this new tool and direction, and to give you reasons to start implementing this as soon as you're on z/OS 1.12. It would be a perfect project for any of your new system programmers because they won't have as many preconceptions of the functions.

z/OSMF 1.11

Figure 1 shows part of the screen after logon. In the left navigation section, you can see the features that are available on z/OSMF 1.11:

- Configuration
 - The Configuration task helps with configuring TCP/IP policies. This is only available from z/OS 1.11 onwards because of the underlying support within TCP/IP 1.11. This facility is based on the free downloadable TCP/IP Configuration Assistant that is currently available for z/OS 1.10.
- Links
 - IBM provides some of the most standard links that a sysprog might use, but you can add any links that you want to.
- Problem Determination
 - The Incident Log facility simplifies the collection and handling of problem management information generated by SVC dumps. This should greatly reduce the effort to locate and manage problems. We go into more detail below.

Figure 1 – Partial Welcome Screen of z/OSMF for z/OS 1.11



- z/OSMF Administration
 - This allows a z/OSMF administrator to add or modify the Links, and allows certain users to access the various functions.

In addition to the links at the left, the functions can also provide:

- Wizards, such as sending a dump to IBM support, to provide step-by-step procedures to perform typical migration or system programming tasks. An example of this is shown in the next section that describes the Incident Log.
- Automated tasks to reduce the time to manually perform some standard functions. This will be a great time saver.

Incident Log

In order to help you see the real benefit of z/OSMF, let's take a look at the Incident Log function of z/OSMF. If a system SVC dump occurs (listed as an ABEND incident) or an operator or authorized program issues an SVC dump (listed as a User Initiated incident), the system can optionally take snapshots of Operlog, Logrec detail, and Logrec summary. Once you turn on snapshots, the default is to collect data for Operlog for 30 minutes prior to the dump, collect Logrec detail for one hour prior to the dump, and collect Logrec summary data for 24 hours prior to the dump. Pretty neat! You can specify which data you want collected in the CEAPRMxx parmlib member. But the default for CEAPRMxx is to turn all snapshots off. It's important to activate snapshots in order to exploit this feature.

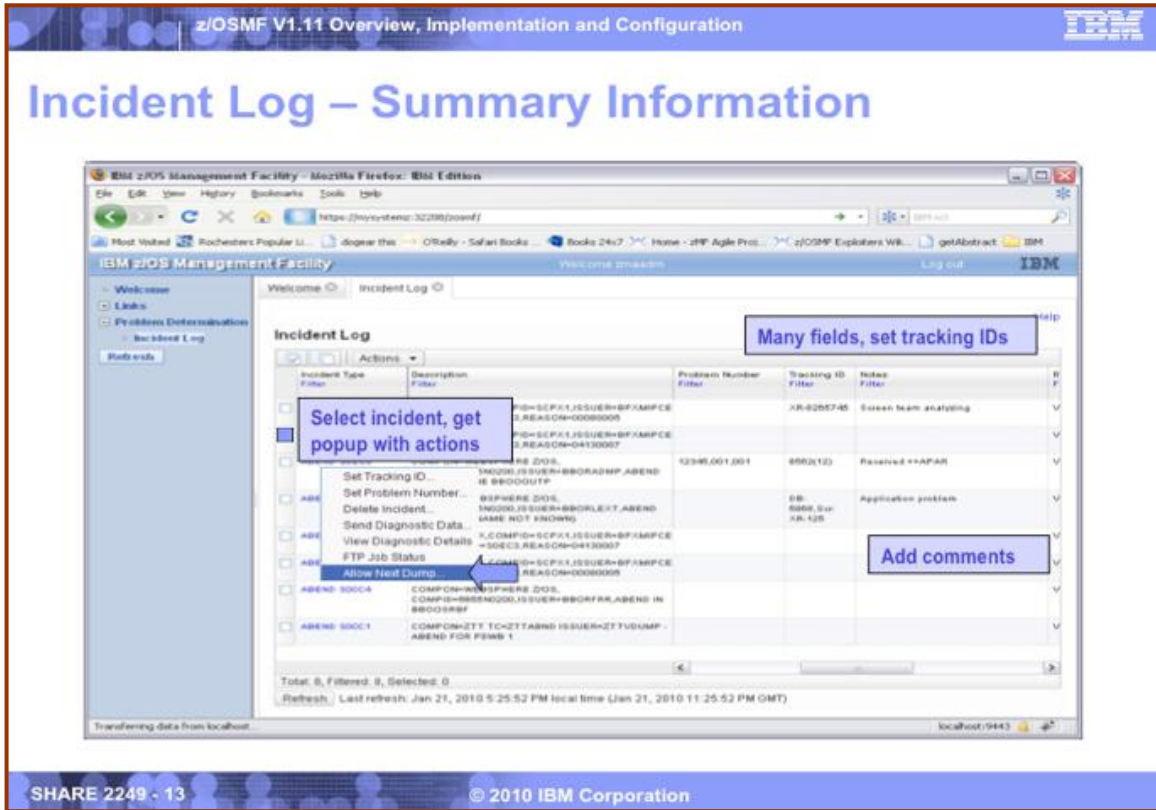
Can you collect and send all of the data needed for a dump to IBM in 30 seconds?
z/OSMF can!

Figure 2 is an example of the Incident Log, so that you can see how easily it can be used. This sysplex-wide directory of all incidents allows one or more sysprogs (or subsystem programmers) to view and work with the dumps. For each dump, you can take several actions, and use wizards that automate many of the actions. As an example, there is a wizard to package up all of the pieces of the diagnostic data collected to send to IBM. IBM claims that this task could take up to 20 minutes when done manually, but can be done with 8 key clicks in 30 seconds using z/OSMF!

As a default, three days of incidents are displayed, although you can display up to 500 incidents. When you select an incident, you can take one of several actions:

- **Set Tracking Id** - Any reference ID you'd like to use. This might be something that corresponds to your problem management system.
- **Set Problem Number** - This could be a PMR, ETR, or ISV problem number.
- **Delete Incident** - Unnecessary events should be deleted in order to keep the list at a minimum, simplify maintenance, and reduce the file sizes.
- **Send Diagnostic Data** - This is the wizard we mentioned that can take all of the diagnostic information, compress it, and FTP it to wherever you want.

Figure 2 - Example of a z/OSMF Incident Log with Actions from Session 2249



- **View Diagnostic Details** - This provides you detailed information about all of the data collected. For example, there might be a primary dump, and some secondary dumps from other systems in the sysplex, and operator logs. z/OSMF collects them, organizes them, and displays them.
- **FTP Job Status** - You can cancel the FTP job, or view/delete the status.
- **Allow Next Dump** - Selecting this updates the DAE (Dump Analysis & Elimination) data set, so that you can capture the next SVC symptom string that might be suppressed by DAE.

Although most sysprogs have developed their own method for collecting and keeping track of all the data for an SVC dump, doesn't this sound a lot easier? And it makes turnover much simpler.

How is it Packaged?

For a variety of reasons, z/OSMF is not integrated into the base z/OS system. It is a separate product that you need to order, but the resulting price is zero. It has a release level that corresponds to the level of z/OS when it was introduced. The first release was z/OSMF 1.11, even though you can install it on z/OS 1.10. The next release was z/OSMF 1.12, and the current release is z/OSMF 1.13.

There are really two parts of the product, with one being referred to as WAS/OEM, and the other referred to as z/OSMF. WAS/OEM is a specially packaged version of WAS V7 for exclusive use by z/OSMF. Even if you've already installed WAS (WebSphere Application Server) on your system for

applications, you will still need to install WAS/OEM for use by z/OSMF. It's quite a different product. See the Announcement Letter in the *References* section for product numbers.

Access is via a Web browser running on Microsoft Windows XP, but not every Web browser. You will need either Mozilla Firefox V2 or V3 (minimum service level 3.0.6) or Microsoft Internet Explorer (IE) V6 or V7. Benchmarks have shown that Firefox delivers better performance than IE. z/OSMF 1.13 provides support for Windows 7 Professional, Firefox 3.5 & 3.6, IE V7 and V8. New browser support is added with each release of z/OSMF.

z/OSMF provides many benefits, but the simplification and automation of system programmer tasks is really impressive. See Figure 3 for an example of this. This example shows the benefits of using the z/OSMF Incident Log.

Figure 3 – Reductions in Sysprog Time Using Incident Log from session 2249

The screenshot shows a slide titled "z/OSMF Problem Determination – Incident log Benefits". It contains a table comparing tasks performed "Without z/OSMF Incident Log" versus "With z/OSMF Incident Log".

	Without z/OSMF Incident Log **	With z/OSMF Incident Log **
Recognizing a system-detected (dumped) problem occurred	Requires 5 to 7 manual steps, plus skill on effective use of IPCS to extract data from each of the dumps. Up to 5-6 minutes	Display in 1 click. Greatly reduced skill required As little as 5 seconds
Allow new dump to be taken for the same symptom	Requires 7 to 12 manual steps, plus skill on effective use of IPCS to locate the dump data set, obtain the symptom string, get into the IPCS DAE display, locate the matching symptom string (could be non-trivial) and indicate TakeNext on the IPCS display Up to 15 minutes	Make the update happen in 3 mouse clicks As little as 10 seconds
Collecting and sending diagnostic data	Requires 7 to 15 manual steps, plus skill to locate the right log files, build and run jobs, rename the output datasets, and use an FTP job to send the different data sets to the target destination. Up to 20 minutes Up to 30 minutes for sysplex components	Send the material in 8 clicks: •Select the incident materials •Specify the FTP destination information •Send the material •Check whether the information was FTP'd successfully As little as 30 seconds

"So easy, even a marketing professional can use it!" – Gita Grube Berg, IBM System z Marketing

** Based on IBM laboratory results, your results may vary

SHARE 2249 - 11 © 2010 IBM Corporation

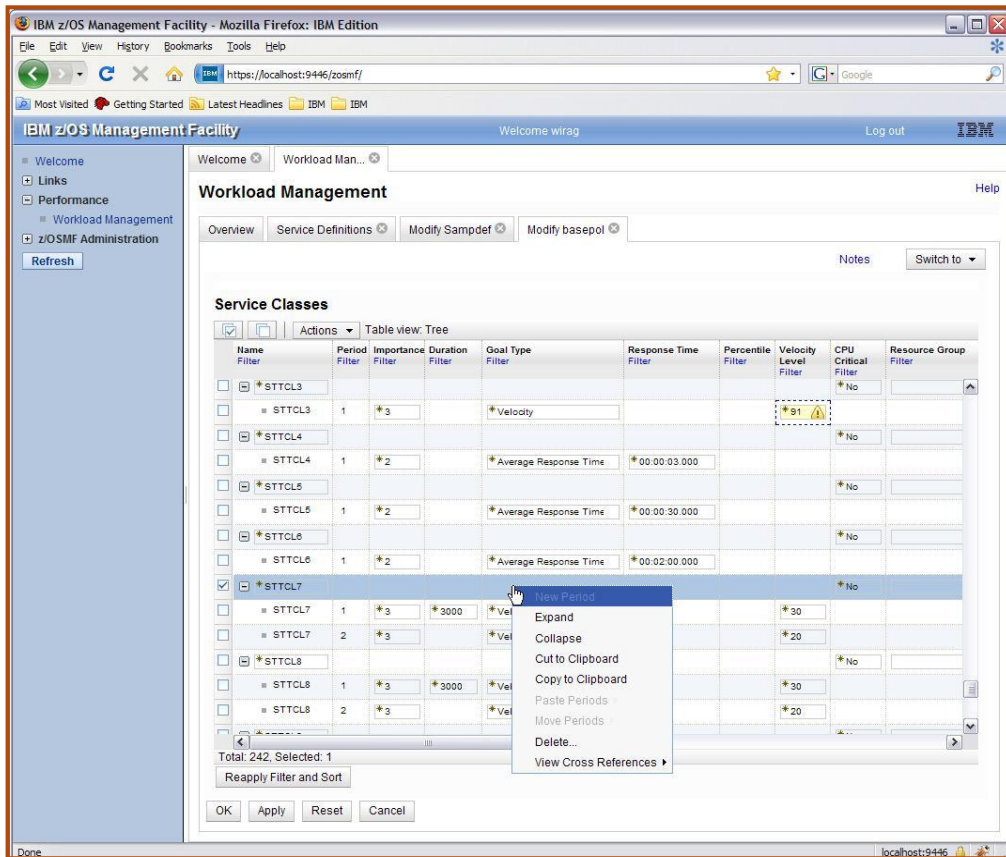
z/OSMF 1.12 Enhancements

The following enhancements were made in z/OSMF 1.12:

- **Performance Menu Item**

This is a new z/OSMF navigation item in z/OSMF 1.12, and contains the Workload Management task, the Monitoring Desktops task, and the Sysplex Status task.

Figure 4 - z/OS WLM Policy Editor - Service Class Periods



■ **Performance Menu - Workload Management**

z/OSMF Workload Management provides four features that should make this worth the time to install z/OSMF. First of all, it will provide the WLM policy in a much easier to use format, (Web-based remember). Not only will you have tables to update instead of using

the current ISPF-based method, but you can print the policies in table format so they're easier to understand.

Second, the new WLM Policy Editor will provide extensive validity checking as you are editing. This is much more than is available today. You can be notified of an error during entry instead of when you try to VALIDATE the policy.

Third, and most important, it will provide recommendations and best practices as you update or create a policy. Wowza! Figure 4 shows you one of the screens from the z/OS 1.12 editor (thanks to the z/OSMF team). This is how you would update the service class periods using the Policy Editor. You can see how much easier this technique is.

And finally, z/OSMF allows you to keep historical copies of your service definition. We thought that this was important enough that we included this support in the very first release of our GoalTender software five years ago. We're glad to see that IBM is now providing this function too.

Users LOVE WLM under z/OSMF!

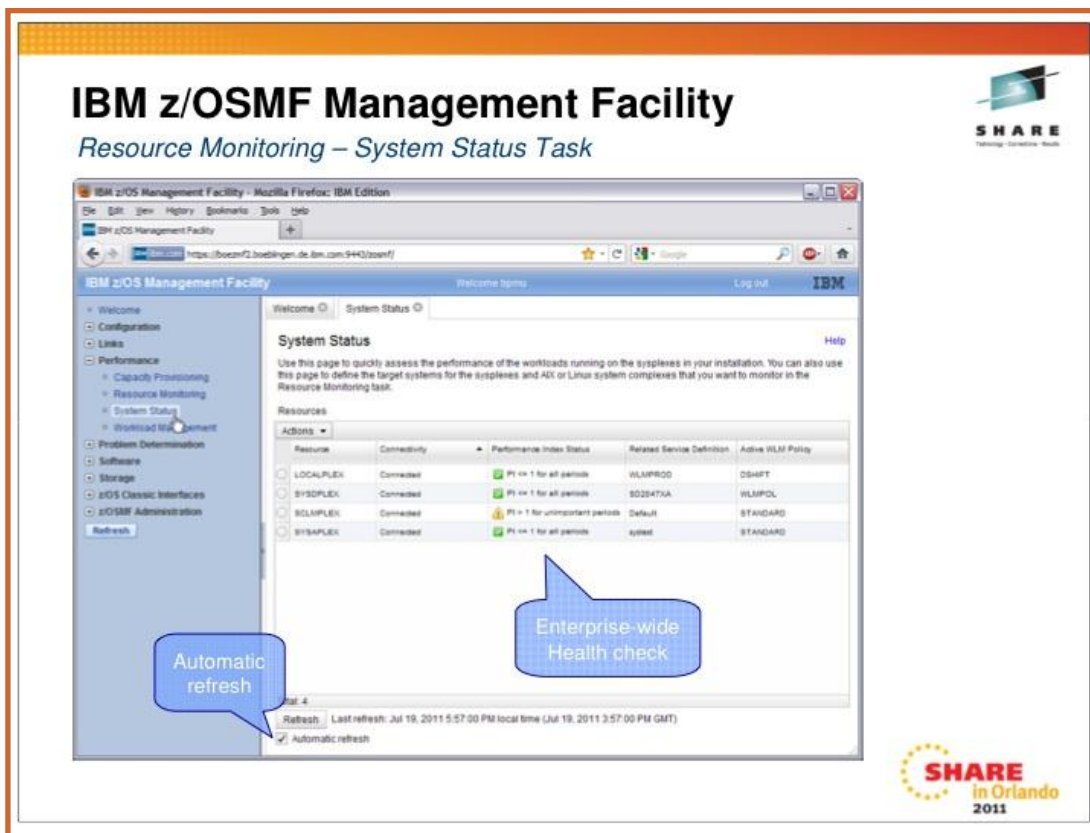
z/OSMF users almost unanimously love the Workload Management feature. They speak very positively about the ease of use, the search feature, the benefit of scrolling through classifications and service classes, the greatly improved report format, and (of course) the warnings provided before they install a policy.

■ **Performance Menu - System Status & Resource Monitoring**

[Note - for much more information, please see **Peter Muench's** (IBM) Orlando SHARE session [10012](#). The figures for the System Status and Resource Monitoring are courtesy of Peter, and come from that handout.]

The Sysplex Status task is provided by RMF, and therefore requires not only RMF, but the RMF Distributed Data Server (DSS). In z/OS 1.13, this has been renamed to 'System Status', so we'll use that term from now on. In a single page, you can see the status of every z/OS sysplex, and every zLinux complex. Figure 5 shows a sample of the System Status page. A red-yellow-green flag shows the health of each complex based on whether service class periods are meeting their goals. If all periods are meeting their goals, the flag is green; if periods with importance levels of 3, 4, or 5 are missing their goals, the flag is yellow; and if periods with important levels of 1 or 2 are missing their goals, the flag is red. Any sysplex or complex running a DDS is available to be included in the display.

Figure 5 – z/OSMF System Status - Courtesy of Peter Muench, © IBM



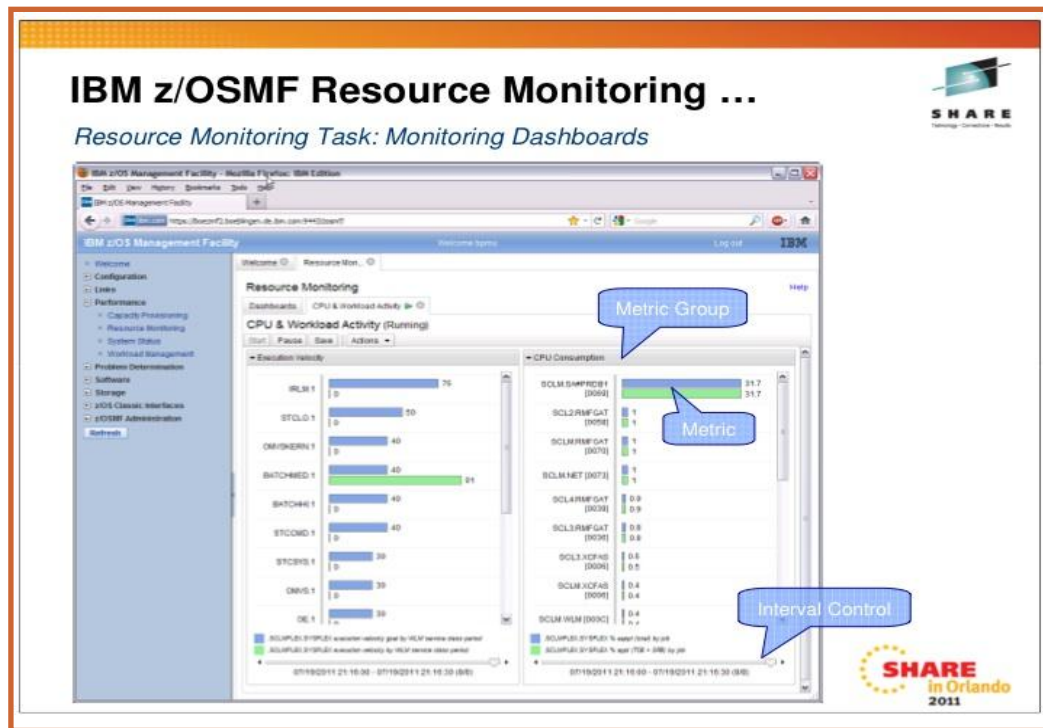
In z/OSMF 1.12, you can download a free tool to gather information for:

- Linux for System z
- Linux for x86 systems

That tool is 'rmfpms', a free Linux data gatherer that can be downloaded from the RMF home page at www.ibm.com/systems/z/os/zos/features/rmf/. It requires the RMF DSS on the host, and an rmfpms task running on each zLinux image. z/OS 1.13 introduced a new component, RMF XP (RMF Cross Platform Performance Monitoring), that is a preferred alternative to the downloadable collector. RMF XP is also available on z/OS 1.12 with APAR [OA36030](http://www.ibm.com/ibm/apar/oa36030) (28Jun2011). In addition to zLinux on a System z machine, RMF XP can collect performance data from AIX on a System p and Linux running on a System x.

If you see a problem on any sysplex or complex, you can use the Monitoring Desktops task (renamed to Resource Monitoring in z/OS 1.13) to drill down and identify the problem. Because the latest documentation uses the term 'Resource Monitoring', we'll use that too. Resource Monitoring provides graphical displays, called dashboards, to show the status of most resources. One example of a dashboard is shown in Figure 6, which shows both execution velocity and CPU consumption by address space.

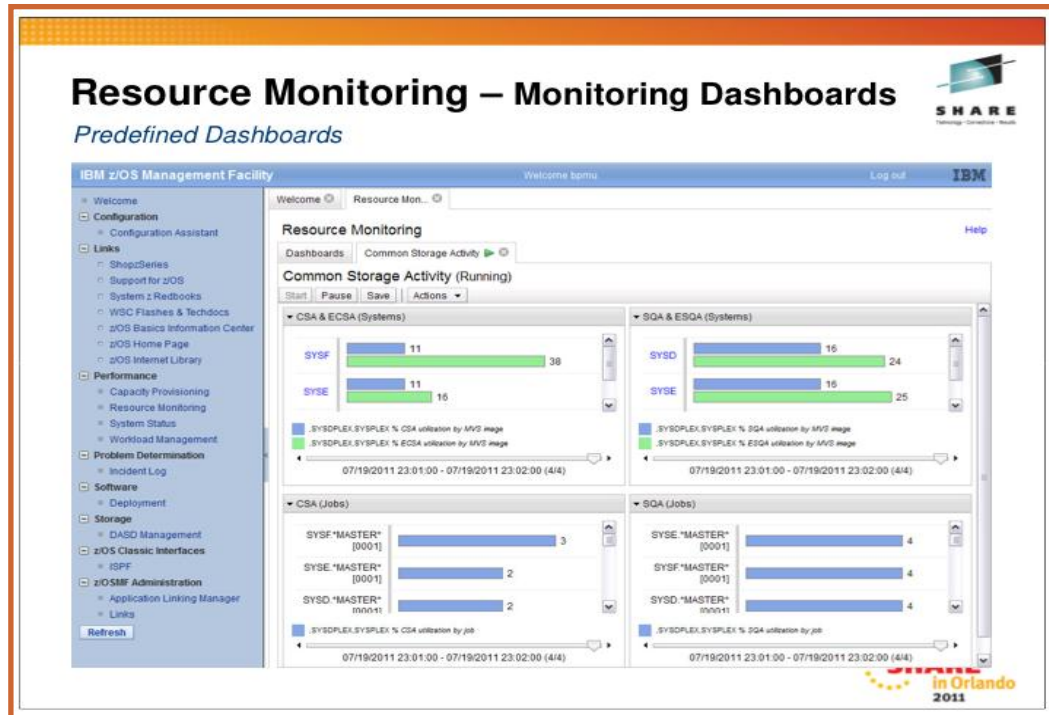
Figure 6 – z/OSMF CPU & Workload Activity - Courtesy of Peter Muench, © IBM



z/OSMF provides several predefined dashboards, although you can create your own. The default display is a snapshot of the current status (e.g. CPU usage or frame usage), but you can change both the interval and the time of the display allowing you to see historically what might have happened. Another example of one of the distributed dashboards is seen in Figure 7, which shows four displays of common storage by systems and jobs. Prerequisites for Resource Monitoring:

1. You need to be consistently collecting data in RMF Monitor III from all z/OS images in a sysplex before z/OSMF can access the data. That doesn't take much of an effort if you're already running RMF, but it may require that you create RMF-owned VSAM data sets for logging if you want to display history in z/OSMF.
2. You need to be running the Distributed Data Server (DSS) address space on (only) one of the z/OS images in a sysplex. Not many installations are doing this, so you're likely to need this step. DDS can be CPU-intensive if you aren't careful about cycle and interval times.

Figure 7 – z/OSMF Storage Monitoring - - Courtesy of Peter Muench, © IBM



■ **Configuration Task**

The Configuration Assistant for TCP/IP added support for several technical enhancements such as IKEv2, new cryptographic algorithms for IPsec and Internet Key Exchange (IKE), configuration of FIPS (Federal Information Processing Standards) 140 cryptographic mode for IKE, certificate trust chains, certificate revocation lists, and enforcement of RFC4301 compliance for IPsec filter rules. There were several usability enhancements as well.

■ **Incident Log Task**

There are several enhancements to Incident Log in z/OSMF 1.12.

- Encryption of the incident files, including dumps.
- Division of the dumps into smaller pieces for transmission in parallel.
- Ability to send additional documentation and attachments with the dump.
- Addition of free form notes for each incident.

- Use of SYSLOG and logrec data sets for a single system in place of the OPERLOG and LOGREC log streams.
- Support for an EMAIL variable in the FTP profiles.

■ **Miscellaneous Items**

- z/OS 1.12 allows you to add any link (even ISVs) to any category, not just the pre-defined z/OS Links categories. Additionally, you can add non-z/OSMF launch points and links to the left navigation bar.
- Several scripts were simplified and changed.
- The Common Information Model (CIM) is no longer part of the z/OSMF configuration process. You will need to configure CIM as a separate step.
- Greatly improved documentation.

z/OSMF 1.13 Enhancements

The following enhancements for z/OSMF 1.13 have been described in manuals and presentations:

■ **Incident Log Enhancements**

The following functions were added to the Incident Log in z/OSMF 1.13:

- The name of the transmitted file is sent in the message.
- The FTP destination and Firewall Proxy settings are shared with the Software Deployment task.
- UNIX system temporary data sets can be used, and can be compressed before sending.
- Provides support for PDUU (Problem Documentation Upload Utility).
- The Retain search string is added to the View Diagnostic Details screen.
- Provides support for [ftp.data](#) when using PDUU.

■ **Performance Enhancements**

As mentioned earlier, 'Monitoring Desktops' has been renamed to 'Resource Monitoring', and 'Sysplex Status' has been renamed to 'System Status'. 'Desktops' have been renamed to 'dashboards'. z/OSMF 1.13 and RMF added support for AIX systems, and began using CIM-based performance data gatherers. The current valid operating systems for z/OS 1.13 are:

- z/OS (using GPMSERVE)
- AIX (using GPM4CIM)
- Linux on System x (using GPM4CIM)
- Linux on System z (using GPM4CIM)
- Linux (using rmfpms, a free Linux data gatherer on the RMF home page at www.ibm.com/systems/z/os/zos/features/rmf/. It requires the RMF DSS on the host.

■ Configuration Assistant Enhancements

- A single instance of the Configuration Assistant can be used to configure both z/OS 1.12 and z/OS 1.13 Communications Servers.
- The Configuration Assistant can retrieve TCP/IP profiles from active stacks in order to provide a list of IP addresses.
- A single policy rule can be defined for multiple stacks.
- Intrusion Detection Services (IDS) adds support for IPv6 and additional attack types, such as out of sequence packet denial of service attacks.

■ z/OSMF Base Enhancements

In z/OSMF 1.13, you have a choice of using your existing Repository mode or switching to SAF mode. IBM's statement of directions indicates that Repository mode will be phased out in a later release. SAF mode allows a tighter integration with z/OS security management: resource names are associated with all z/OSMF tasks and links; resource class profiles manage authorization to z/OSMF managed resources; and SAF groups are used to represent z/OSMF roles. As an example, this can be used to set different authorities for installing a WLM service policy, modifying one, and viewing one.

z/OS Software
deployment –
is installing
z/OS next?

A new programmatic interface for z/OS called a RESTful web service (also called REST API) is an interface to z/OS to allow web service applications (such as javascript, PHP, Perl, etc.) to submit batch jobs to JES2 and JES3. APIs exist to allow the applications to manage the jobs after submission (get status of a job, get list of spool files for a job, get list of jobs, cancel a job, change the job class for a job, and purge a job).

Application linking is a third new feature of base z/OSMF. This allows you to define links between applications. One of the examples that's given is to let you define an event, such as 'LookAt', and then define an 'event handler' (and action) to will link to the event. Another example is to allowing linking an Incident Log entry to SDSF to look at a job log.

■ Software Deployment

This new function of z/OSMF 1.13 is really the start of a future method of installing z/OS (in my opinion). The purpose of software deployment is to "make cloning of installed software simpler and safer". This includes all SMP/E installed software, both IBM and non-IBM. One of the first steps to using software deployment is to start implementing IBM recommended best practices. If you think about it, there are almost as many methods of installing and upgrading z/OS as there are installations. IBM's technique here is to start with the outer ring first - the remote systems that are simply clones of the base system. Once all of those start to look alike, it might then be possible to create a single method of installing and upgrading that base system. (At least that's what I surmise.) But for now, the z/OSMF Software Deployment provides the following:

- Verification of cross system and cross product software prerequisites.
- Verification that software fixes aren't regressed.
- Cloning all parts of the software.
- Cloning the inventory (SMP/E CSI) with the software.
- Support for local cloning (single system or system-to-system within a sysplex) or remote cloning (system-to-system across a network and multiple sysplexes).

I am in LOVE
with the z/OSMF
ISPF interface!

The Software Deployment task can be found under a new *Software* tab on the left navigation bar. This is IBM's recommended path for cloning and deployment.

■ **Capacity Provisioning**

Currently, a Capacity Provisioning Control Center (CPCC) is used to administer and monitor capacity provisioning policies that are managed by the Capacity Provisioning Manager (CPM). CPM is used to control capacity on demand. CPCC is available as a separate Windows-based stand-alone client. z/OSMF 1.13 provides a Capacity Provisioning task to perform part of the CPCC monitoring functions. This task manages CIM connections to CPMs and provides reports for three areas - domain status, active configuration, and active policy. This task is found under *Performance* on the left navigation bar.

Figure 8 – Partial Welcome Screen for z/OSMF 1.13



■ **z/OSMF ISPF Interface**


This is my third most favorite part of z/OSMF. In case you're curious, my favorite is the Incident Log, and my second favorite is the Workload Management task. A new item on the left navigation bar is *z/OS Classic Interfaces*, where you'll find the ISPF task. With this task, you can have web-based ISPF panels that are scrollable, but still support standard PF keys. You can define up to four panes on a single monitor, or use tabs so that multiple panes are available. Popups notify you of TSO messages that come in. Figure 9 is from **Anuja's** Orlando SHARE session [9804](#) and shows two methods of displaying ISPF screens.

So that you can appreciate the improvements between z/OSMF 1.11 and 1.13, compare the new z/OS 1.13 Welcome screen in Figure 8 with the original screen shown in Figure 1 on page 3.

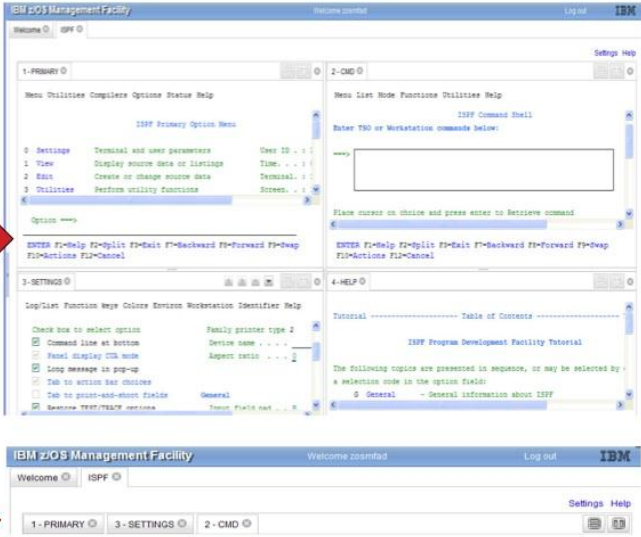
Figure 9 – z/OSMF 1.13 ISPF Support – Extract from Anuja's SHARE session 9804 © IBM


z/OSMF and ISPF (R13)

Work with existing interfaces



- **Enables system programmer to perform tasks from one interface.**
 - Also makes ISPF applications URL Web-accessible for linking and launching from other applications
- Up to 4 panes
 - Panes can be sized. Each pane can have multiple ISPF sessions, tabs can be moved between panes
 - Shows ISPF Menu bar, Command line, Function keys
- Also available in tabbed format





Resource Requirements

In May of 2011, IBM updated a white paper, [WP101779](#) (10May2011), and an APAR, [PM34596](#) (9May2011), that better described the resource requirements of z/OSMF. The requirements were stated as "2 GB of central storage" and "available CPU resource equivalent to a processor with a processor capacity index (PCI) of at least 120." Many people read this to mean that the CEC have at least 2 GB and 120 MIPS. But the real meaning was that the LPAR running z/OSMF needs to provide 2 GB of real storage to z/OSMF itself, and that you need a minimum of 120 MIPS just for z/OSMF. This naturally caused a great deal of concern among those of us who thought that z/OSMF would be a great solution for small installations. We described this situation, the paper, and the feedback in our [Cheryl's List #151](#) on August 26, 2011.

Yes, you need
2 GB of memory
and 120 MIPS
(but understand when)

The original paper is still on the website and labeled as 'z/OSMF Resource Requirements for z/OS 1.12', but an additional paper is added and labeled as 'z/OSMF Resource Requirements for z/OS 1.13'. Here are some major points that we saw in the updated z/OSMF 1.13 document:

- I think that the most important change in the 1.13 paper is a sentence that qualifies the 120 MIPS requirement as being needed "during periods of high CPU usages for certain scenarios." And they state that it may be possible to run z/OSMF with less than these requirements (including the 2 GB of storage), depending on what you're doing. The CPU-intensive tasks (specifically Software Deployment and WLM editing of a very large policy) are listed below. The paper recommends that you NOT use Software Deployment in a configuration less than the recommended 120 MIPS and 2 GB.
- The paper also points out more strongly that a large portion of z/OSMF is written in Java, and can therefore take advantage of a zAAP if available (or a zIIP if using zAAP on zIIP).
- Lab results were run on a z196-701 machine with a single dedicated CPU (about 1200 MIPS) and 8 GB of memory. Some testing was also done using 1/10th of the CPU (120 MIPS) and 2 GB of memory. Steady state z/OSMF (no activity) took about .5% of the 120 MIPS, or about 6 MIPS. zAAP or zIIP eligibility ranged between 4% and 95%. The testing configuration had about 150 incidents for Incident Log, two SMP/E target zones (each with 631 data sets), and both a large (1 MB) and a small (35 K) WLM policy.
- Relative CPU usage per browser click for a variety of requests ranged from smallest to largest:

- ISPF display job status (i.e. took the least amount of CPU time)
- ISPF browse
- HCD display devices
- Capacity Provisioning configuration
- Capacity Provisioning domain
- Incident Log display of 30 incidents
- WLM edit of a small policy
- Incident Log view detail
- Configuration Assistant

Incident Log display of 150 incidents

Incident allow dump

WLM edit of a large policy (this took twice as much CPU than the Incident allow dump) *[Note: in the z/OSMF 1.12 paper, they noted that IWMARINO (the ISPF WLM) took even more resources than z/OSMF.]*

Software Deployment took so long in comparison to even WLM editing a large policy, that they needed a new chart for scaling. But the CPU time depends very much on how many clones and how many FMIDs are involved.

- A great addition to the paper is a graph showing which tasks can use the zAAPs, with WLM editing a large policy being able to move about half of the CPU time to a zAAP. Most of the Configuration Assistant could move to a zAAP, and half of the ISPF browse could use a zAAP.
- z/OSMF startup time took about one minute when running on the 1200 MIPS CPU, and took about 12 minutes on the 120 MIPS configuration.
- The paper also included several recommendations on reducing startup time.

Here are a few points from the z/OSMF 1.12 paper that we didn't mention previously:

- Resource monitoring with three concurrent desktops (i.e. views) required less than 1.5% of a z10 CPU (or about 13 MIPS).
- Eligible zAAP usage ranged between 25% and 75% of the total CPU time.

So, what do you need to run z/OSMF? Just what the White paper says - a minimum of 2 GB of storage dedicated to z/OSMF, and a minimum of 120 MIPS **when doing z/OSMF functions that are CPU-intensive**. The addition of a zAAP or a zIIP can reduce those CPU requirements. You can somewhat minimize the impact of z/OSMF on your system by attempting to use the longer functions at off-peak CPU times. Also realize that your response times on a test LPAR will be very different from those on your production system.

I've talked to several z/OSMF customers who are running z/OSMF in much smaller environments. While it wasn't lightening fast, it was certainly acceptable. Anyone installing z/OSMF for the first time should be aware that their first experience with z/OSMF will probably be on a small test LPAR that doesn't have many resources. So performance may suffer.

As an example of a steady state (no z/OSMF activity), **Mary Anne Matyaz** and **Mark Brown** working at US Customs found that on their development LPAR (400 MIPS, 10 GB), z/OSMF took about .26 MIPS, and startup took only .13 MIPS. The working set size of all z/OSMF components (including WAS OEM) ranged from 1 GB to 1.5 GB. Mary Anne's SHARE session on z/OSMF Experiences can be found here - share.confex.com/share/117/webprogram/Session9737.html. In a later email Mary Anne sent me the results of additional testing. She did a lot of work under z/OSMF, including working with a fairly large WLM policy, and the z/OSMF STCs took less than 10 MIPS. Not part of z/OSMF, but just for comparison, the FTP of a dump to IBM took about 160 MIPS.

What does this mean to large sites? Almost nothing. Except that if you need to use all that 120 MIPS during peak processing, your rolling 4-hour average might increase, and thus the cost of all other software may increase. Of course, you can always consider putting z/OSMF in its own LPAR.

Or you could avoid performing the tasks that use a lot of CPU until a less busy time of day. (E.g. you could update your huge WLM policy or work with a large incident log at 8 am instead of 11 am.) [Note: if you are able to move some of the z/OSMF work to a zIIP or zAAP, you might be able to avoid increasing your rolling 4-hour average.]

But small sites, and anyone testing in a small development LPAR, might encounter some performance issues if they do not provide enough resources. One site found that z/OSMF startup took over 20 minutes until they added more storage. The white paper indicates that z/OSMF startup should take about 2 minutes on a z10.

Implementation Recommendations

Here are our primary recommendations for getting started with z/OSMF:

1. If you have a small machine, then you will probably need to wait for z/OSMF. I fully expect IBM to reduce the resource requirements at some point in the future. If you're willing to wait a long time for the initialization to complete, and you don't mind if z/OSMF abends now and then, then you can try it to get a feel for it. But don't bother with production until either IBM comes to your rescue, or you get a bigger machine.
2. Order the product along with your base z/OS system because it's not automatic, even if there is no charge for the product. The product number for z/OSMF is 5655-S28 (and Subscription and Service is 5644-S29). Ensure that you have an additional 5 GB of DASD space for downloading the WAS OEM product.
3. Accept that there may be some problems along the way, but realize that IBM wants to make this easier for you. If you run into **ANY** problems, be sure to contact the z/OSMF support team for help. (Obviously, you should first be following the steps in the installation manuals.)
4. Review the *IBM Hot Topics* articles, the z/OSMF Redbook, the installation manuals, the SHARE presentations, and the white paper on Resource Requirements (see *References* on page 18). For installation on z/OS 1.12, I recommend starting with **Greg Daynes'** (IBM) Anaheim SHARE session [8698](#) (z/OSMF 1.12 Implementation). For installation on z/OS 1.13, I recommend starting with **Anuja Deedwaniya's** (IBM) Orlando SHARE session [9733](#) (z/OSMF 1.13 Implementation).
5. Take the time to implement the prerequisites for each task before you even begin the install for z/OSMF. These are IBM recommended best practices, so they provide a benefit even before you get to z/OSMF. These include: SYSLOG, system logger, OPERLOG, CIM, CEA, DAE, dynamic dump data sets, and RMF DSS. The SHARE handouts in the previous item also provide a list of the prerequisites for each of the tasks.

Have patience
installing z/OSMF;
ask for help

6. Wait until z/OS 1.12 or 1.13 to install z/OSMF. I cannot recommend installing the z/OSMF 1.11 release due to the complexity of the install.
7. Ensure that you have either enough resources (e.g. 2 GB of memory and 120 MIPS) on your test LPAR or are willing to let WAS timeout and ABEND z/OSMF.
8. Install Firefox on your workstation. IBM's benchmarks have shown that the performance of Firefox exceeds that of Microsoft's Internet Explorer for the z/OSMF client.
9. Start with the easiest tasks first. I would start with adding some links, adding the Configuration Assistant (if your TCP/IP administrator wants to use it), and then implementing Workload Management. If you are installing for the first time in z/OS 1.13, then try the ISPF function first. This might be a good time to run through the *Hands On Lab* from SHARE session [9735](#) to become familiar with z/OSMF. (Bypass the Incident Log exercises until you have that installed.)
10. Once you're familiar with z/OSMF, then you can consider installing the Incident Log task and the RMF Resource Management task. If you have the prerequisites in place, both of these tasks are easy to install, but you may want to take time to tailor them to your needs. You may also need to set up naming standards and new procedures for anyone dealing with dumps. This may require some training for the other administrators, such as the DB2 or CICS admins.
11. If you are running zLinux LPARs or zLinux guests under z/VM, then install the RMF collector for zLinux so that the System Status panels can show both z/OS and zLinux systems.
12. Please let us know about your experiences, and we'll share them with our other readers. We're especially interested in any experiences with the new z/OSMF 1.13 functions.

Latest References

There are several new references since we last discussed z/OSMF, so here's a complete list as of the end of September 2011.

IBM Publications & Websites

To get a general understanding of z/OSMF, I suggest reading this entire article. Then read the *IBM Hot Topics* issues listed below. The articles are short and show you why each facility is important. For implementation, I would start with the SHARE session 9733 (below), which tells you which manuals to use. If you're not implementing z/OSMF right away, the SHARE presentations in the next section from **Anuja Deedwaniya** provide more information about each of the tasks.

- [GA32-0631-04](#) - *IBM WAS OEM Edition for z/OS Configuration Guide*.
- [GI11-2886-00](#) - *z/OSMF V1R11 Program Directory*
- [GI11-2886-01](#) - *z/OSMF V1R12 Program Directory*
- [GI11-2886-02](#) - *z/OSMF V1R13 Program Directory*

- IBM Education Assistant - [Training for z/OSMF](http://pub-lib.boulder.ibm.com/infocenter/ieduasst/stqv1r0/index.jsp?topic=/com.ibm.iea.zosmf/plugin_coverpage.html) - pub-lib.boulder.ibm.com/infocenter/ieduasst/stqv1r0/index.jsp?topic=/com.ibm.iea.zosmf/plugin_coverpage.html
- [IBM Hot Topics Issue 21](#) (Aug2009) - This contains many articles about z/OSMF, including an introduction, incident log, setting up operlog and logrec, CEA, and system logger for z/OSMF.
- [IBM Hot Topics Issue 23](#) (Aug2010) - This issue contains several more articles about z/OSMF, including an introduction to z/OSMF, a new look for z/OSMF, using WLM and the resource monitoring facility of z/OSMF, enhancements to the Incident Log, and sizing the system logger.
- [IBM Hot Topics Issue 25](#) (Aug2011) - This issue contains additional articles about z/OSMF, including two articles on software deployment, RMF XP, and z/OSMF 1.13 enhancements.
- IBM z/OSMF website - www.ibm.com/systems/z/os/zos/zosmf/
- [SA38-0652-01](#) - *z/OSMF V1R11 User's Guide* (without APAR [PK97274](#))
- [SA38-0652-02](#) - *z/OSMF V1R11 User's Guide* (for APAR [PK97274](#))
- [SA38-0652-03](#) - *z/OSMF V1R12 Configuration Guide* (notice the name change from above)
- [SA38-0652-04](#) - *z/OSMF V1R12 Configuration Guide* (for APARs [OA29534](#), [OA35023](#), [PM24215](#), and [PM24843](#))
- [SA38-0652-05](#) - *z/OSMF V1R13 Configuration Guide*
- [SA38-0654-00](#) - *z/OSMF Messages*. This first version starts with z/OSMF 1.13, but is useful for earlier releases.
- [SC33-7990-18](#) - *z/OS V1R12 RMF User's Guide*
- [SC33-7990-19](#) - *z/OS V1R13 RMF User's Guide*
- [SG24-7851-00](#) - *z/OS Management Facility*. The Redbook draft is based on z/OSMF 1.12, and was installed on a system that had most of the prerequisites and best practices in place. Unfortunately, that doesn't correspond to real life in most installations. It's a great book for getting z/OSMF installed and exploited. (22Jul2011, Updated 5Aug2011)
- [T1012012](#) - *Supported z/OSMF V1R11 Functionality, by z/OS Release*. This ATS FAQ Technote showed which z/OSMF functions were available for z/OS 1.10 and 1.11. (7Apr2011)
- [T1012682](#) - *z/OSMF Incident Data Management*. This ATS Technote (FAQ) describes how to best manage incidents stored by the Incident Log function. (7Apr2011)
- [WP101779](#) - *z/OS V1.12 Management Facility Resource Requirements*. See the discussion under *Resource Requirements* on page 15. (Updated 5Oct2011)

SHARE Presentations

The latest SHARE conference in Atlanta (March 2012) had the following presentations (www.share.org):

- [10635](#) - **Anuja Deedwaniya** - *What's New in z/OS 1.13*
- [10652](#) - **Greg Daynes** - *z/OS Software Deployment Application*
- [10638](#) - **Skip Robinson & Micheline Bissell** - *z/OSMF 1.13 User Experience*

There were some excellent SHARE presentations on z/OSMF in Orlando (August 2011)

(www.share.org):

- [9732](#) - **Kurt Quackenbush** - *z/OS Software Deployment* (this was repeated in Atlanta, March 2012, session [10652](#))
- [9734](#) - **Anuja Deedwaniya** - *z/OSMF Overview*
- [9733](#) - **Anuja Deedwaniya** - *z/OSMF 1.13 Implementation & Configuration*
- [9735](#) - **Anuja Deedwaniya** - *z/OSMF Hands-On Lab*. Even if you weren't there for the lab, the instructions are very handy for learning how to use z/OSMF once you get it up and running.
- [9737](#) - **Mary Anne Matyaz & Brad Carson** - *z/OSMF User Experience*
- [9804](#) - **Anuja Deedwaniya** - *What's New in z/OSMF V1.13?*
- [10002](#) - **Glenn Anderson** - *zEnterprise eXposed! Part 1: The Intersection of WLM, RMF, and z/OSMF Performance Management*
- [10011](#) - **Peter Muench** - *RMF: The Latest and Greatest*
- [10012](#) - **Peter Muench** - *Manage Your Workloads & Performance with z/OSMF*

There were also some excellent SHARE presentations on z/OSMF in Anaheim (February 2011)

that are useful (www.share.org):

- [8512](#) - **Skip Robinson & Miguel Mercado** - *z/OS 1.12 User Experiences*
- [8696](#) - **Karla Arndt & Anuja Deedwaniya** - *z/OS Problem Determination Update: z/OSMF Incident Log, Runtime Diagnostics, PFA, and New Technologies*
- [8698](#) - **Greg Daynes** - *z/OSMF 1.12 Implementation & Configuration*
- [8757](#) - **Mary Anne Matyaz & Brad Carson** - *z/OSMF User Experiences*
- [9061](#) - **Glenn Anderson** - *Understanding WebSphere App Server OEM for z/OSMF Sys-progs*. Glenn provides an excellent description of WAS OEM, as well as recommendations on how to set up goals for WAS OEM in WLM.

z/OSMF APARs

Before installing or implementing z/OSMF, be sure to check out the corresponding PSP bucket, such as Upgrade ZOSMF113. You can also do a search of z/OSMF in your APAR database. ■